

Board oversight of AI transformation

August 2026



Table of contents

Introduction	3
01 Govern AI like a transformation, not a tech initiative	5
02 Align on where to lead, lag, or exit	8
03 Steer the talent and culture needed to unlock AI	12
04 Guide the shift to a workforce of people and AI agents	15
05 Oversee the risks and controls that let AI scale	18
06 Monitor outcomes and risks	21
Conclusion	24

Introduction

Since the launch of ChatGPT in late 2022, AI has rapidly evolved from a breakthrough technology into a core driver of business transformation. The pace of change has been rapid, and companies are now moving beyond experimentation to redesign how work gets done across products and services, decision-making, operations, and the workforce itself. As with other major technology shifts, the long-term impact will depend less on the tools themselves and more on how efficiently companies redesign the business around them.

A company's success with AI will likely depend on clear strategic choices, disciplined investment, workforce and operating model change, and the ability to scale adoption for competitive advantage. It will also depend on building trust through effective risk management, strong governance, and responsible use.



As technology advances and becomes more accessible, companies face dynamic and evolving transformation. AI strategies may need to adapt over time, requiring ongoing reassessment of priorities and investments. History shows that we often misjudge technological change, overestimating short-term breakthroughs while overlooking its deeper, long-term impact. AI is no exception.

The board's role is to oversee whether management is making the right decisions, building the right conditions for success, and adapting with enough agility as the technology, competitive landscape, and regulatory expectations continue to evolve. Directors, on behalf of shareholders, have to govern AI in the short term while also challenging management on the longer-term choices that will shape business strategy.

Six key actions boards can take to provide effective oversight of AI transformation

- | | | | |
|----------|--|----------|--|
| 1 | Govern AI like a transformation, not a tech initiative | 4 | Guide the shift to a workforce of people and AI agents |
| 2 | Align on where to lead, lag, or exit | 5 | Oversee the risks and controls that let AI scale |
| 3 | Steer the talent and culture needed to unlock AI | 6 | Monitor outcomes and risks |

Govern AI like a transformation, not a tech initiative

AI governance starts with clear accountability

As AI moves from experimentation to enterprise transformation, the board's role is to oversee the strategic decisions that drive the transformation and hold management accountable for results. That means understanding how AI is changing the company's strategy, operating model, workforce, corporate culture, and risk profile.

Clarify board and committee oversight. Typically, the full board oversees the strategy for AI transformation because of the direct connection to business strategy, enterprise transformation, and long-term shareholder value. However, effective AI oversight should not sit just with the board or a single committee. AI is an enterprise-wide transformation that affects strategy, operations, talent, risk, financial reporting, compliance, and other areas. As a result, oversight responsibilities will cut across the full board and its committees.

Some companies may have a board-level technology committee that oversees the transformation effort. These committees are relatively rare, though slowly increasing in number. Only 17% of S&P 500 boards have a technology committee, but this is up from just 15% in 2021.¹

1 ESGAUGE, 2026.

While the full board or a technology committee can take the lead on oversight, each board committee will have oversight responsibility for different aspects of the company's AI transformation and risk profile. For example:



Audit committee: AI risks in the ERM process; risk oversight allocation and AI model governance; and risks and opportunities associated with the use of AI in financial reporting, audit plans, and compliance.



Compensation committee: The impact of AI on the overall talent strategy and workforce, compensation plans, incentives, and corporate culture; and risks and opportunities related to the use of AI in the human capital function.



Nominating/governance committee: Communications around the AI story to investors and stakeholders; board composition; use of AI in board oversight practices; and continuing education and upskilling on AI.

Establish clear management accountability. Some companies have a chief AI officer (CAIO) who coordinates AI transformation. Other companies may have the chief strategy officer or chief operating officer lead this effort. What matters most is not the title but that there is accountability for a strategy-driven, outcomes-focused AI plan. This individual will lead the AI agenda and coordinate across the business with business unit, finance, legal, risk, human resources, and other C-suite leaders. Some companies are also establishing management-level councils with key leaders to align priorities, build a risk and governance structure, and drive AI transformation.

Build AI fluency in the boardroom. Given the speed of change and impact of AI, boards should possess—or have access to—the skills and expertise needed to oversee this area effectively. Notably, 71% of directors say AI is the board capability most in need of strengthening.²

Boards can assess whether they need additional education, outside expertise, or different director profiles to keep pace. Adding a director with specialized AI knowledge is one option to consider based on the company's business model, industry, and risk profile. While directors are not expected to be AI experts, they do need sufficient fluency to oversee AI with confidence. Directors can build that fluency through direct engagement with AI tools, board education sessions led by management or outside specialists, board-level AI certification programs, external programs, and even the company's own employee AI training.

Using AI in the boardroom is becoming part of effective oversight

Board adoption of AI remains in its early stages, with only 40% of directors reporting they use AI in their oversight role, up modestly from 35% last year.³ Boards can also use AI to strengthen their own work such as using technology as a strategic thought partner, synthesizing board materials, conducting research and scenario modeling, and preparing more effectively and efficiently for meetings. The board's use of AI should support better decision-making, not replace it. Confidentiality and legal risks should be understood and managed.

To responsibly integrate AI into board processes and workflows, the board should consider adopting a proactive governance framework. To learn more, read **Using AI in the boardroom—new opportunities and challenges**.


² PwC, 2026 Annual Corporate Directors Survey, forthcoming.

³ PwC, 2026 Annual Corporate Directors Survey, forthcoming.


Align on where to lead, lag, or exit

Not every AI opportunity deserves leadership

As AI investment increases, the board's role is to oversee how AI is driving the business strategy and long-term shareholder value. AI is not a strategy in itself, but an accelerator of business strategy. The board should challenge management on whether the company is making the right strategic choices and moving bold enough and fast enough. This should be a dialogue and alignment with management about where the company is in its AI journey, what is realistic for the company and sector, and where the greatest business value is likely to emerge.



of CEOs report either revenue growth or cost reductions from their AI initiatives



say their companies are 'stuck' and have realized neither revenue nor cost benefits

Source: PwC's 29th Global CEO Survey

Not every company is seeing returns from AI yet. Some are delivering revenue gains and cost savings, while many others remain stuck in experiments. What separates companies that achieve ROI is not more activity, but focused investment aligned with strategic priorities rather than scattered experiments with little enterprise impact.

Align on strategic priorities. Boards should expect to hear from management on how the industry is evolving, where disruptors may emerge, and whether the company's AI posture remains competitive over time. Directors should push management to make explicit choices about where the company wants to lead with AI initiatives, where it intends to keep pace, and where it may choose to exit or deprioritize with a focus on business outcomes.

- **Lead.** Leading initiatives' focus on a small number of high-value bets where AI can truly differentiate the company by improving customer experience, workforce productivity, operating performance, or business model advantage.
- **Lag.** Lagging is a conscious choice. It can be appropriate where the business case is less compelling, the company is not yet ready to scale, or the market is still maturing, and waiting may allow for adoption of a more proven solution.
- **Exit.** Exiting can free up capital, management attention, and organizational capacity to focus on the areas that matter most. As AI changes the economics of certain activities, some areas may no longer justify the investment required to remain competitive.

The company's choices should reflect the maturity of the AI technology. It is important to remember that AI has never seen a use case it doesn't like, but that doesn't mean that every use case will produce meaningful outcomes. Some AI capabilities are mature and ready to scale, while others are still emerging and may require greater investment and experimentation before they deliver value. Boards should push management to pick its spots and match investment to where AI can realistically create value.

Back AI use with the right building blocks

The strongest AI strategies where companies choose to lead are grounded in five core enablers:

- Make sure AI investments power the business strategy
- Pick areas where AI capabilities are proven to deliver value
- Embed AI into end-to-end workflows and not just layer onto existing complexity
- Deploy or source the best talent to make the transformation a success
- Address readiness of data and technology capabilities

Treat data and technology as strategic assets. Data readiness remains a gating factor for many AI transformations, and even the best strategies can stall when data quality, usage rights, consent, security, and governance are not sufficiently mature. Poor data practices can limit the value of AI initiatives and increase operational, legal, privacy, and reputational risks.

Boards should understand major technology decisions about AI models, cloud and compute infrastructure, proprietary versus third-party tools, and data architecture because these choices affect speed, cost, risk management, vendor dependence, and resilience in both the short- and long-term. Many companies are seeking to preserve flexibility and optionality in their AI systems, including avoiding overdependence on any single model or vendor as the market continues to evolve.

Centralize AI investment decisions. Boards should oversee AI capital allocation like any other major investment with clear priorities, centralized processes, and measured outcomes. AI spend goes well beyond tools and models; it can include data readiness, computing capacity, technology upgrades, external partners, and employee training and upskilling, along with recurring costs for cloud, computing, and model usage. Many AI costs are consumption-based rather than fixed. Boards should understand how management monitors these ongoing operating costs and whether they remain aligned with expected business value.

Questions boards can ask management


- How is AI being used by our competitors? Are there disruptors for our industry?
- Are we making the right bets that truly move our business model, not just funding experiments?
- Are we treating AI like an investment with centralized funding and measuring ROI over time?
- Do we have the technology infrastructure and high-quality data necessary for our focused AI transformation investments?
- How do we build agility into our AI use with the rapidly evolving landscape, considering investments and vendor optionality?

S&P 500 benchmarking points to a widening AI performance gap

New benchmarking is beginning to give boards a more practical lens on AI performance. Looking across S&P 500 companies from 2022 to 2025, excluding AI foundation builders such as Alphabet, Microsoft, and NVIDIA, companies that invested more than 0.5% of revenue in AI outperformed their sector median total shareholder return by 21%, while those investing less underperformed by 2%. The message for boards is not simply to spend more—it is that companies treating AI as an enterprise-wide growth and transformation lever appear to be pulling ahead of those still relying on isolated pilots and low-risk automation.

The same pattern appears in operating performance. Companies investing more than 0.5% of revenue in AI outperformed their sector median revenue growth by 3% and EBITDA growth by 7%, while lower-investing peers were roughly flat relative to sector medians.

To learn more about AI benchmarking, read [**Agentic AI reinvention: Making AI agents accretive to the P&L.**](#)



Note: Based on research data across S&P 500 companies from 2022 to 2025, excluding foundation builders, e.g. NVIDIA, Microsoft, Alphabet
Source: S&P, CapIQ, Earnings Reports, PwC Analysis

Steer the talent and culture needed to unlock AI

Great talent drives AI success

AI transformation is as much about people as it is about technology. While AI capabilities are becoming more accessible, the ability to create lasting business value increasingly depends on having the right leaders, skills, and culture. The board's role is to oversee whether management has the talent strategy, organizational capabilities, and change management needed to embed AI across the enterprise.

Boards should expect to hear more frequently from the chief human resources officer (CHRO), alongside business and technology leaders. Discussions should extend beyond workforce planning to include leadership capability, critical AI skills, workforce readiness, and how the organization fosters a culture that encourages learning, responsible experimentation, and continuous adaptation.

Advance the workforce capabilities needed for AI. Technology alone rarely determines the success of AI initiatives. The companies creating the greatest value from AI are those with the right people leading, designing, governing, testing, and adopting AI-enabled ways of working. Competitive advantage will increasingly come from human capability—the ability to rethink work, make key decisions, manage risk, and drive adoption at scale.

Boards will want to discuss with management how the organization is developing leaders, attracting and retaining critical AI talent, and equipping the broader workforce with the skills needed to work effectively in an AI-enabled environment.

The adoption gap is real

54% of workers say they have used AI in the last 12 months, but



only 14% say they use GenAI daily at work



That gap suggests many organizations still have work to do to turn awareness into embedded behavior change


Source: PwC's Global Workforce Hopes and Fears Survey 2025

Foster a culture where AI adoption sticks. AI transformation will not succeed without employee trust. Employees are more likely to embrace AI when they trust leadership, understand what is changing, and feel supported in learning new ways of working.

Boards should expect management to have a change management program that includes leadership alignment, targeted upskilling and reskilling, transparent communication, and clear accountability for adoption. Leadership should also address common barriers to adoption, including change fatigue, lack of familiarity with AI, concerns about job displacement, and uncertainty about evolving roles.

Measure outcomes, not just activity. Measuring AI adoption and its connection to meaningful business outcomes can be challenging. Boards should avoid relying solely on activity-based metrics, such as sessions, users, or tokens consumed, as these may reflect experimentation, or even misuse, rather than value creation. Instead, companies should pair usage data with more outcome-oriented indicators, including weighing employee engagement and workforce sentiment measures with progress against strategic goals. Taken together, these measures can help boards assess whether AI transformation is becoming embedded in how work gets done and whether it is improving workforce effectiveness and enterprise performance.

Many workers view AI's impact on their jobs positively



Q: To what extent, if at all, do you feel the following emotions about how AI may affect your work?

(Showing only 'To a very large extent' and 'To a large extent' answers)

Source: PwC's Global Workforce Hopes and Fears Survey 2025

Questions boards can ask management

- Do we have the leadership, talent, and workforce capabilities needed to execute our AI strategy?
- What is our change management program to drive employee adoption and use?
- How are we measuring and monitoring AI adoption, workforce readiness, and employee sentiment?
- Are leaders and incentives aligned to reward learning and responsible experimentation and innovation?


Guide the shift to a workforce of people and AI agents

AI agents execute work—people remain accountable

As companies embed AI into core workflows and move to a workforce that includes both people and AI agents, existing talent and operating models will be disrupted. The most successful companies treat agentic AI adoption as a workforce and operating-model transformation, not a technology rollout. This requires redesigning work and not simply automating tasks.

Agents can handle routine inquiries, reason through tasks, execute multi-step processes, generate code, and extend the reach of knowledge workers. But the real value does not come from the technology alone—it comes from a human-led, tech-powered model in which people guide, oversee, and collaborate with agents. It comes from combining AI capabilities with strong human judgment. Boards should expect management to rethink end-to-end workflows, team structures, decision rights, and accountability, rather than simply layering AI onto existing complexity.

AI agents and multi-agent workflows



Agent: An AI agent is a system given a high-level objective or mission. It can reason through a task, use tools such as application programming interfaces (APIs), databases, search, or code environments, and adjust its actions based on feedback until the objective is completed.



Multi-agent workflow: A multi-agent workflow is a setup in which multiple AI agents, each with a dedicated objective, work together in the same environment to complete a broader task.

Keep human judgment at the center. As more execution work shifts to agents, human roles will move toward judgment, exception handling, and risk oversight with humans remaining at the center of accountability for decisions and outcomes. Boards should expect management to define where human judgment is non-negotiable, particularly in higher-impact decisions. Management should determine when a human must approve, review, override, or remain accountable for outcomes, and whether AI outputs are explainable and auditable enough for the context.

Rethink work, not just the workforce. Directors should expect management to explain how the operating model and talent strategy will evolve to support a workforce of people and AI agents. Management should demonstrate how work will be redesigned including decision rights, accountability, governance, team structures, and critical skills as AI agents become embedded in core business processes.

Companies should also reassess their sourcing model. Management may need to revisit what should be brought in-house and what should remain or be outsourced based on talent and AI capabilities.

An example of how AI agents are reshaping customer service work

In a traditional customer service model, most work is handled by people across a tiered support structure. Frontline representatives answer routine questions, more experienced staff handle exceptions and complaints, and managers oversee performance and escalation.

In an agent-enabled workforce, AI agents could take on a meaningful share of the routine work including answering common questions, summarizing customer history, drafting responses, processing simple requests, and routing more complex issues to the right person. Human employees then shift toward higher-value work such as resolving exceptions, handling sensitive situations, improving service quality, and stepping in when the agent is wrong or uncertain.

That changes the skills the workforce needs. Some employees may help design, train, test, and improve the agents, working with workflows, prompts, data, and controls. At the same time, experienced or upskilled customer service professionals become even more important because they understand customer needs, service standards, escalation points, and what good judgment looks like in practice. Their role evolves from mainly doing the work to supervising the work, solving issues, and helping improve how humans and agents work together.

For boards, the key point is that this is not just automation. It is a redesign of roles, skills, and accountability across the service model.

Questions boards can ask management

- How are we preparing for a workforce that includes both people and AI agents?
- Are we embedding AI into our core processes and operating model, or simply layering it onto existing complexity?
- What will our workforce look like in three years in size, skills, insourcing and outsourcing, and geographic footprint, and what is the transition path from today?

Oversee the risks and controls that let AI scale

Trust enables AI to scale

Transforming with AI while maintaining trust is not a separate compliance exercise. It is the enabling foundation that allows AI to scale safely and at speed. For directors, the challenge is to make sure AI risks are understood, prioritized, and managed, while not slowing innovation. Governance over risks should help the company move faster with confidence and build trust with stakeholders.

Oversee the foundation for trusted AI. Boards should start by confirming that management has established a clear AI risk management foundation, including policies, a current inventory of where AI is being used across the enterprise, a common risk taxonomy, and baseline controls. That inventory should include AI developed internally, AI embedded in third-party software, and AI tools used by employees.

AI risks should be integrated into the company's broader risk management approach, with a common language for assessing exposure across risks, including model, data, infrastructure, and user risks. There are also legal and compliance and process impact risks that need to be managed. Human oversight remains critical for AI technology, and management should be clear about who owns AI risk decisions, models and agents, escalation procedures, and remediation. Emerging AI risks, especially given the pace of technological advancement, can exceed what current risk environments were designed to manage, and boards should expect management to update risk structures accordingly.

Controls to manage risks should be designed into AI initiatives from the outset. Boards should expect embedded controls across AI systems during the design, testing, deployment, and post-deployment phases.

Key AI risks

Model risks



Risk related to the training, development, and performance of the AI system itself, including conceptual soundness, reliability of output, and oversight

Data risks



Risk related to the collection, processing, storage, management, and usage of data during the training and operation of the AI system

Infrastructure risks



Risk related to the acquisition, implementation, and operation of an AI system in a broader software and technology environment

User risks



Risk related to the intentional or unintentional misuse, manipulation of, or attacks against an AI system

Legal and compliance risks

Noncompliance risks with applicable laws, rules, and regulations

Process impact risks

Unforeseen or unmitigated risks that arise from integrating AI into an existing process or workflow without proper assessment

Focus on risk-tiering and high-risk models. A single governance process for every AI use case creates bottlenecks. Directors will want to discuss with management how governance is differentiated by AI risk tier, with streamlined processes for lower-risk uses and more intensive review, controls, escalation procedures, and oversight for higher-risk systems. Clear criteria should be used to determine risk tiers, including factors such as sensitive data use, external exposure, autonomous action, business impact, and regulatory sensitivity.

Boards should receive additional reporting on the highest-risk models or use cases that create the greatest exposure. This should include AI systems that make significant judgments or decisions or that could have a high consequence if they fail, drift, or are misused. For these models, it may be appropriate to have independent validation or assurance of the model by internal audit or a third party.

Have significant AI risks escalated. Management should be moving to continuous monitoring of AI-enabled workflows to identify incidents, control breakdowns, and other risks. Continuous monitoring is necessary as AI systems can drift, behave unpredictably, generate harmful or misleading outputs, or deteriorate over time. Boards should be informed of major incidents early, with clear escalation triggers and resilience and response readiness playbooks updated for a potential AI incident.

The evolving regulatory landscape

The regulatory environment for AI remains fragmented. In the United States, the federal posture is still broadly pro-innovation, but there is no single comprehensive federal AI law. Instead, companies face a mix of agency guidance, existing laws, and state-level requirements. On March 20, 2026, the White House released a National Policy Framework for AI that purports to preempt state-level AI regulations to avoid a patchwork of various AI laws. Outside the United States, many global regulatory initiatives are underway, and the landscape continues to evolve.

Boards should ask management how it is actively monitoring relevant regulatory developments and building governance that can adapt as requirements evolve.

Questions boards can ask management

- Where are we most exposed today across our AI ecosystem?
- Are we differentiating governance by AI risk tier or applying similar oversight across most use cases in practice?
- Which high-risk AI models or use cases should the board review more closely, and why?
- What triggers are in place to escalate an AI incident to the board before the exposure becomes significant?

Monitor outcomes and risks

Business outcomes matter more than AI activity

The board's role is to monitor whether management's AI use, investments, and operating changes are producing business value, whether capital is being allocated with discipline over time, and where risks are emerging. Boards should expect regular reporting on the AI transformation journey, including progress against goals and milestones, challenges encountered, risk management, and whether strategic pivots may be needed.

As AI transformation scales, boards need reporting that goes beyond activity updates to show what is delivering value, where risks are emerging, and whether the company is building the conditions for sustained success. Many companies are still in the early stages of using AI to transform and won't yet have a holistic set of metrics upon which to base results. But as the company's use of AI expands, these metrics become more important.

AI reporting falls short

Companies are still in the early stages of measuring AI value. Few directors report that they receive quality information that links AI investments to business performance.



Reinforce management accountability. Boards should understand who owns AI strategy, execution, and governance, and how often key leaders will engage with the board. It can be helpful for the CAIO or CIO and business or functional leaders to present together to confirm alignment between technology efforts and business outcomes.

Measure what matters. Measuring AI ROI can be challenging, and every company will have its own criteria for success. But directors should still expect clear and relevant metrics tied to the goals of the transformation. Depending on the business and industry, those metrics may include efficiency gains, cost savings, new or incremental revenue, enhanced quality, cycle-time improvements, and better customer experience. Metrics such as AI usage, model count, and tool adoption may not be as impactful measures for ROI.

Build stakeholder trust through transparency. Clear communication about AI use, governance, and risk mitigation helps build confidence internally and externally. Directors should understand how management demonstrates responsible AI use and fosters trust with investors, regulators, employees, and other stakeholders. They should also pay attention to messaging and oversee whether it is evidence-based, avoids overstatement, and guards against AI washing. Boards should be prepared to answer investor questions about how the board oversees AI transformation, whether AI initiatives are tied to business strategy and the outcomes, where oversight sits, and whether the board has the right skills and expertise.

Questions boards can ask management

- What are our next steps based on the strategic outcomes, model performance, business outcomes, and other metrics reported?
- What is the most significant issue you have dealt with since we last met, and how was it resolved?
- How should our AI use change based on the successes and learnings to date?

Board reporting to support effective AI oversight

Key AI transformation reporting

- Strategic plan for key AI initiatives and tie(s) to overall business strategy
- Execution roadmap, budget, and related KPIs and progress
- Talent strategy planning
- Risk foundation (policies, AI inventory, risk taxonomy, third-party risk management)
- Major relevant regulatory developments
- Benchmarking against industry peers

Illustrative AI transformation metrics to monitor progress

Strategic plan

- AI investment as a percentage of revenue and by function
- Percentage of priority AI use cases deployed at scale
- ROI/value realization for key strategic business cases (e.g. market share, cost reduction, process time, AI customer interaction)
- Number of processes and roles redesigned for human-plus-agent workflows
- Number of milestones completed on time and delayed

Risk oversight

- Percentage of AI use cases risk-tiered
- Percentage of high-risk systems passing validation/testing and independently assessed
- Number of significant AI incidents/exceptions escalated to senior leadership
- Percentage of significant AI incident issues remediated within target timeframe

Training and culture

- Percentage of priority AI-related roles filled versus approved plan
- Percentage of employees completing AI training within required timeframe
- Percentage of eligible employees using approved AI tools daily (or alternate adoption metric)
- Percentage of employees who feel confident in AI use, experience value, and trust AI adoption

Conclusion

The value-generating potential of AI is significant, as are its risks. But companies likely won't realize its transformative effects without focused investment and disciplined execution. Boards have an important role in helping management approach AI as an enterprise transformation and providing oversight of the decisions that matter most.





How PwC can help

To have a deeper discussion about how this topic might impact your business, please contact your engagement partner or one of the following PwC contacts.

Ray Garcia

Partner, Governance Insights Center Leader
ray.r.garcia@pwc.com

Dan Priest

Chief AI Officer
dan.priest@pwc.com

Jenn Kosar

AI Assurance Leader
jennifer.kosar@pwc.com

Barbara Berlin

Managing Director, Governance Insights Center
barbara.berlin@pwc.com